


Application for United States Letters Patent

for

**PROCESSING TASKS WITH FAILURE RECOVERY**

by

**Dale E. Gulick**

CERTIFICATE OF MAILING UNDER 37 C.F.R. § 1.10	
EXPRESS MAIL NO.:	EL 656 271 840 US
DATE OF DEPOSIT:	JANUARY 11, 2002
<p>I hereby certify that this paper or fee is being deposited with the United States Postal Service "EXPRESS MAIL POST OFFICE TO ADDRESSEE" service under 37 C.F.R. 1.10 on the date indicated above and is addressed to: Box Patent Application, Assistant Commissioner for Patents, Washington, D.C. 20231.</p> <p> Signature</p>	

## PROCESSING TASKS WITH FAILURE RECOVERY

### BACKGROUND OF THE INVENTION

#### **1. FIELD OF THE INVENTION**

This invention relates generally to computing systems, and, more particularly, to processing tasks with failure recovery in a computer system.

#### **2. DESCRIPTION OF THE RELATED ART**

Fig. 1A illustrates an exemplary computer system 100. The computer system 100 includes a processor 102, a north bridge 104, memory 106, Advanced Graphics Port (AGP) device 108, a network interface card (NIC) 109, a Peripheral Component Interconnect (PCI) bus 110, a PCI connector 111, a south bridge 112, a battery 113, an AT Attachment (ATA) interface 114 (more commonly known as an Integrated Drive Electronics (IDE) interface), an SMBus 115, a universal serial bus (USB) interface 116, a Low Pin Count (LPC) bus 118, an input/output controller chip (SuperI/O<sup>TM</sup>) 120, and BIOS memory 122. It is noted that the north bridge 104 and the south bridge 112 may include only a single chip or a plurality of chips, leading to the collective term "chipset." It is also noted that other buses, devices, and/or subsystems may be included in the computer system 100 as desired, e.g. caches, modems, parallel or serial interfaces, SCSI interfaces, etc.

The processor 102 is coupled to the north bridge 104. The north bridge 104 provides an interface between the processor 102, the memory 106, the AGP device 108, and the PCI bus 110. The south bridge 112 provides an interface between the PCI bus 110 and the peripherals, devices, and subsystems coupled to the IDE interface 114, the SMBus 115, the

USB interface 116, and the LPC bus 118. The battery 113 is shown coupled to the south bridge 112. The Super I/O™ chip 120 is coupled to the LPC bus 118.

The north bridge 104 provides communications access between and/or among the 5 processor 102, memory 106, the AGP device 108, devices coupled to the PCI bus 110, and devices and subsystems coupled to the south bridge 112. Typically, removable peripheral devices are inserted into PCI “slots,” shown here as the PCI connector 111, that connect to the PCI bus 110 to couple to the computer system 100. Alternatively, devices located on a 10 motherboard may be directly connected to the PCI bus 110. The SMBus 115 may be “integrated” with the PCI bus 110 by using pins in the PCI connector 111 for a portion of the SMBus 115 connections.

The south bridge 112 provides an interface between the PCI bus 110 and various 15 devices and subsystems, such as a modem, a printer, keyboard, mouse, etc., which are generally coupled to the computer system 100 through the LPC bus 118, or one of its predecessors, such as an X-bus or an Industry Standard Architecture (ISA) bus. The south bridge 112 includes logic used to interface the devices to the rest of computer system 100 through the IDE interface 114, the USB interface 116, and the LPC bus 118. The south bridge 112 also includes the logic to interface with devices through the SMBus 115, an 20 extension of the two-wire inter-IC bus protocol.

Fig. 1B illustrates certain aspects of the south bridge 112, including reserve power by the battery 113, so-called “being inside the RTC (real time clock) battery well” 125. The south bridge 112 includes south bridge (SB) RAM 126 and a clock circuit 128, both inside 25 the RTC battery well 125. The SB RAM 126 includes CMOS RAM 126A and RTC RAM

126B. The RTC RAM 126B includes clock data 129 and checksum data 127. The south bridge 112 also includes, outside the RTC battery well 125, a CPU interface 132, power and system management units 133, and various bus interface logic circuits 134.

5 Time and date data from the clock circuit 128 are stored as the clock data 129 in the RTC RAM 126B. The checksum data 127 in the RTC RAM 126B may be calculated based on the CMOS RAM 126A data and stored by BIOS during the boot process, such as is described below, e.g. block 148, with respect to Fig. 2. The CPU interface 132 may include interrupt signal controllers and processor signal controllers.

10 Fig. 1C illustrates a prior art remote management configuration for the computer system 100. A motherboard 101 provides structural and base electrical support for the south bridge 112, the PCI bus 110, the PCI connector 111, the SMBus 115, and sensors 103A and 103B. The NIC 109, a removable add-in card, couples to the motherboard 101, the PCI bus 15 110, and the SMBus 115 through the PCI connector 111. The NIC 109 includes an Ethernet controller 105 and an ASF microcontroller 107. The Ethernet controller 105 communicates with a remote management server 90, passing management data and commands between the ASF microcontroller 107 and the remote management server 90. The remote management server 90 is external to the computer system 100.

20

An industry standard specification, generally referred to as the Alert Standard Format (ASF) Specification, defines one approach to “system manageability” using the remote management server 90. The ASF Specification defines remote control and alerting interfaces capable of operating when an operating system of a client system, such as the computer 25 system 100, is not functioning. Generally, the remote management server 90 is configured to

monitor and control one or more client systems. Typical operations of the ASF alerting interfaces include transmitting alert messages from a client to the remote management server 90, sending remote control commands from the remote management server 90 to the client(s) and responses from the client(s) to the remote management server 90, determining and 5 transmitting to the remote management server 90 the client-specific configurations and assets, and configuring and controlling the client(s) by interacting with the operating system(s) of the client(s). In addition, the remote management server 90 communicates with the ASF NIC 109 and the client(s)' ASF NIC 109 communicates with local client sensors 103 and the local client host processor.

10

When the client has an ACPI-aware operating system functioning, configuration software for the ASF NIC 109 runs during a "one good boot" to store certain ASF, ACPI, and client configuration data.

15

The transmission protocol in ASF for sending alerts from the client to the remote management server 90 is the Platform Event Trap (PET). A PET frame consists of a plurality of fields, including GUID (globally unique identifier), sequence number, time, source of PET frame at the client, event type code, event level, sensor device that caused the alert, event data, and ID fields.

20

Many events may cause an alert to be sent. The events may include temperature value over or under a set-point, voltage value over or under a set-point, fan actual or predicted failure, fan speed over or under a set-point, and physical computer system intrusion. System operation errors may also be alerts, such as memory errors, data device errors, data controller 25 errors, CPU electrical characteristic mis-matches, etc. Alerts may also correspond to BIOS or

firmware progression during booting or initialization of any part of the client. Operating system (OS) events may also generate alerts, such as OS boot failure or OS timeouts. The ASF Specification provides for a “heartbeat” alert with a programmable period typically one minute but not to exceed 10 minutes, when the client does not send out the heartbeat, or “I am 5 still here,” message.

Client control functions are implemented through a remote management and control protocol (RMCP) that is a user datagram protocol (UDP) based protocol. RMCP is used when the client is not running the operating system. RMCP packets are exchanged during reset, power-up, and power-down cycles, each having a different message type. The remote management server 90 determines the ASF-RMCP capabilities of the client(s) by a handshake protocol using a presence-ping-request that is acknowledged by the client(s) and followed-up with a presence-pong that indicates the ASF version being used. The remote management server 90 then sends a request to the client to indicate the configuration of the client, which the client acknowledges and follows with a message giving the configuration of the client as stored in non-volatile memory during the “one good boot.” The RMCP packets include a contents field, a type field, an offset field, and a value field. 15

RMCP message transactions involve a request from the remote management server 20 90, a timed wait for an acknowledgement followed by a second timed wait for a response. If either of the time limits for the acknowledgement or the response is exceeded, then the remote management server 90 knows that either the client needs some of the packets resent or the client has lost contact due to failure of either the client or the communications link.

The ASF NIC 109 must be able to report its IP (Internet protocol) address (or equivalent) without the intervention of the operating system. Thus, the ASF NIC 109 should be able to receive and reply to ARP (Address Resolution Protocol) requests with the operating system, not interfere with ARP packets when the operating system is running, and 5 wake-up for ARP packets when configured to do so. Note that ACPI includes waking-up for ARP packets as a standard configuration.

The following information is sent to the remote management server 90 from the client as an indication of the configuration of the client: an ACPI description table identifying sensors and their characteristics, ASF capabilities and system type for PET messages, and the client's support for RCMP and the last RCMP command; how the client configures an optional operating system boot hang failure recovery timer; and the SMBIOS identification of the UUID/GUID for PET messages. ASF objects follow the ASL (ASF Software Language) naming convention of ACPI.

15

Based in part on the above-described features, modern computer systems are becoming more and more robust than their predecessors. Computer systems today process a fairly large number of tasks at any given time. As the number of tasks that are processed increases, the likelihood that some of these tasks may not successfully complete (because of errors, for example) also increases. Errant or hung tasks, for example, may adversely affect the performance of the computer system. As such, recovery from these failed tasks is desirable.

25

## **SUMMARY OF THE INVENTION**

In one aspect of the present invention, a method is provided for processing tasks with failure recovery. The method includes storing one or more tasks in a queue, wherein each task has an associated exit routine, and determining at least one task to process based on a priority scheme. The method further includes processing the at least one task, and calling the exit routine based on determining that the task has not completed processing within a preselected period of time.

15  
THE JOURNAL OF

In another aspect of the invention, an apparatus is provided for processing tasks with failure recovery. The apparatus comprises a queue and a controller. The queue includes a task picker stored therein. The controller, which is communicatively coupled to the queue, is adapted to determine if at least one task other than the task picker is stored in the queue and execute the task based on determining that at least one task other than the task picker is stored in the queue. The controller is further adapted to execute the task picker in response to executing the task and to continue to execute the task picker until a preselected event occurs.

#### BRIEF DESCRIPTION OF THE DRAWINGS

The invention may be understood by reference to the following description taken in conjunction with the accompanying drawings, in which like reference numerals identify similar elements, and in which:

Fig. 1A illustrates a block diagram of a prior art computer system, Fig. 1B illustrates a block diagram of a prior art south bridge, and Fig. 1C illustrates a prior art remote management arrangement;

Figs. 2A and 2B illustrate block diagrams of embodiments of computer systems having remote management arrangements, according to various aspects of the present invention;

5 Fig. 3 illustrates a block diagram of an embodiment of an ASF south bridge including integrated ASF, ACPI, and/or Ethernet capabilities, according to various aspects of the present invention;

10 Fig. 4 illustrates a block diagram of an embodiment of the ASF south bridge including ASF registers in the RTC battery well of the ASF south bridge, according to various aspects of the present invention;

15 Fig. 6 illustrates a flowchart an embodiment of a method for booting a computer system including the ASF south bridge of Fig. 4, according to one aspect of the present invention; and

20 Figs. 7A and 7B illustrate flowcharts of embodiments of method for operating a computer system including the ASF south bridge of Fig. 4, according to various aspects of the present invention;

25 Fig. 7C illustrates a block diagram of a polling engine that may be employed in the computer system of Figs. 3A and 3B, according to various aspects of the present invention;

Fig. 7D illustrates an exemplary address table that may be employed by the polling engine of Fig. 7C, according to various aspects of the present invention;

Figure 8 illustrates a block diagram of a south bridge that may be employed in the computer system of Figs. 3A and 3B, according to various aspects of the present invention;

5       Figure 9 illustrates a flow diagram of a master control loop that may be employed in the computer system of Figs. 3A and 3B, according to various aspects of the present invention

10      Figure 10 depicts a flow diagram of an interrupt service routine that may be employed with the master control loop of Figure 9, according to various aspects of the present invention

15      Figure 11 illustrates a flow diagram of an alternative embodiment of the master control loop of Figure 9, according to various aspects of the present invention; and

20      Figure 12 depicts a flow diagram of an interrupt service routine that may be employed with the master control loop of Figure 11, according to various aspects of the present invention.

25      While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof have been shown by way of example in the drawings and are herein described in detail. It should be understood, however, that the description herein of specific embodiments is not intended to limit the invention to the particular forms disclosed, but on the contrary, the intention is to cover all modifications, equivalents, and alternatives falling within the spirit and scope of the invention as defined by the appended claims.

**DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS**

Illustrative embodiments of the invention are described below. In the interest of clarity, not all features of an actual implementation are described in this specification. It will, of course, be appreciated that in the development of any such actual embodiment, numerous 5 implementation-specific decisions must be made to achieve the developers' specific goals, such as compliance with system-related and business-related constraints, which will vary from one implementation to another. Moreover, it will be appreciated that such a development effort might be complex and time-consuming, but would nevertheless be a routine undertaking for those of ordinary skill in the art having the benefit of this disclosure. The use of a letter in association with a reference number is intended to show alternative 10 embodiments or examples of the item to which the reference number is connected.

The following co-pending U.S. Patent Applications are hereby incorporated by reference in their entireties, as if set forth fully herein:

15 [LPC Extension Application] "Method And Apparatus For Extending Legacy Computer Systems", U.S. Patent Application No. 09/544,858, filed on April 7, 2000, whose inventor is Dale E. Gulick; and

[Secure Execution Mode Applications] U.S. Patent Application No. 09/852,372, entitled, "Secure Execution Box and Method," filed on May 10, 2001, whose inventors are 20 Dale E. Gulick and Geoffrey S. Strongin;

U.S. Patent Application No. 09/852,942, entitled, "Computer System Architecture for Enhanced Security and Manageability," filed on May 10, 2001, whose inventors are Geoffrey S. Strongin and Dale E. Gulick;

U.S. Patent Application No. 09/853,395, entitled, "Enhanced Security and Manageability using Secure Storage in a Personal Computer System," filed on May 11, 2001, whose inventors are Geoffrey S. Strongin and Dale E. Gulick;

U.S. Patent Application No. 09/853,446, entitled, "Resource Sequester Mechanism," filed on 5 May 11, 2001, whose inventor is and Dale E. Gulick;

U.S. Patent Application No. 09/853,447, entitled, "Integrated Circuit for Security and Manageability," filed on May 11, 2001, whose inventors are Dale E. Gulick and Geoffrey S. Strongin;

U.S. Patent Application No. 09/853,225, entitled, "System Management Mode Duration and Management," filed on May 11, 2001, whose inventors are Geoffrey S. Strongin and Dale E. Gulick;

U.S. Patent Application No. 09/853,226, entitled, "Mechanism for Closing Back Door Access Mechanisms in Personal Computer Systems," filed on May 11, 2001, whose inventor is Geoffrey S. Strongin;

15 U.S. Patent Application No. 09/854,040, entitled, "Cryptographic Randomness Register for Computer System Security," filed on May 11, 2001, whose inventor is Dale E. Gulick;

U.S. Patent Application No. 09/853,465, entitled, "Cryptographic Command-Response Access to a Memory in a Personal Computer System," filed on May 11, 2001, whose 20 inventor is Geoffrey S. Strongin;

U.S. Patent Application No. 09/853,443, entitled, "Protection Mechanism for Biometric Input Data," filed on May 11, 2001, whose inventors are Dale E. Gulick and Geoffrey S. Strongin;

U.S. Patent Application No. 09/853,437, entitled, "Personal Computer Security Mechanism," 25 filed on May 11, 2001, whose inventors are Geoffrey S. Strongin and Dale E. Gulick;

U.S. Patent Application No. 09/853,335, entitled, "Asset Sharing between Host Processor and Security Hardware," filed on May 11, 2001, whose inventors are Geoffrey S. Strongin and Dale E. Gulick;

5 U.S. Patent Application No. 09/853,234, entitled, "Interruptible and Re-enterable System Management Mode Programming Code," filed on May 11, 2001, whose inventors are Geoffrey S. Strongin and Dale E. Gulick;

U.S. Patent Application No. 09/871,084, entitled, "Locking Mechanism Override and Disable for Personal Computer ROM Access Protection," filed on May 30, 2001, whose inventors are Frederick D. Weber and Dale E. Gulick;

10 U.S. Patent Application No. 09/871,511, entitled, "Monotonic Counter Mechanism for Computer System Security," filed on May 30, 2001, whose inventors Frederick D. Weber and Dale E. Gulick;

15 U.S. Patent Application No. 09/870,890, entitled, "Secure Booting of a Personal Computer System," filed on May 30, 2001, whose inventors are Geoffrey S. Strongin, Dale E. Gulick, and Frederick Weber; and

U.S. Patent Application No. 09/870,889, entitled, "External Locking Mechanism for Personal Computer Memory Locations, filed on May 30, 2001, whose inventors are Geoffrey S. Strongin, Dale E. Gulick, and Frederick Weber.

20 The following non-patent documents are hereby incorporated by reference in their entirety, without prejudice and without disclaimer, as if set forth fully herein:

[ASF] Alert Standard Format Specification, 1.03, 20 June 2001, DSP0114, and earlier version, at <http://www.dmtf.org/spec/asf.html>;

[ACPI] *Advanced Configuration and Power Interface Specification*, 2.0, 27 July 2000, and earlier version, <http://www.teleport.com/~acpi/spec.htm>;

[RFC1157] *A Simple Network Management Protocol*, <http://www.ietf.org/rfc/rfc1157.txt>;

[CIM] *CIM Standards*, <http://www.dmtf.org/spec/cims.html>;

[IPMI] *Intelligent Platform Management Interface Specification v1.0, rev 1.1*, August 26, 1999, and earlier versions, <http://developer.intel.com/design/servers/ipmi/>;

5 [RFC1188] *IP and ARP on FDDI Networks*, <http://www.ietf.org/rfc/rfc1180.txt>;

[FRU] *IPMI Field Replaceable Unit (FRU) Information Storage Definition*, v1.0, 16 September 1998, and earlier versions,  
<ftp://download.intel.com/design/servers/ipmi/fru1010.pdf>;

[MTLS] *Metolious ACPI/Manageability Specification*, v1.0, 30 April 1999,  
<http://developer.intel.com/ial/metolious/index.htm>;

[NDCPM] *Network Device Class Power Management Reference Specification*, v1.0a, 21 November 1997, <http://www.microsoft.com/hwdev/specs/PMref/PMnetwork.htm>;

[PET] *Platform Event Trap Specification*, v1.0, 7 December 1998, and earlier versions,  
<ftp://download.intel.com/design/servers/ipmi/pet100.pdf>;

15 [SCMIS] *SMBus Control Method Interface Specification*, v1.0, 10 December 1999, and earlier versions, <http://www.smbus.org/specs/index.html>;

[SMBIOS] *System Management BIOS Reference Specification*, v2.3.1, 16 March 1999, and earlier versions, <ftp://download.intel.com/ial/wfm/smbios.pdf>;

[SMBUS\_2.0] *System Management Bus (SMBus) Specification*, v2.0, 03 August 2000, and  
20 earlier versions, <http://www.smbus.org/specs/index.html>; and

[RFC\_UDP] *User Datagram Protocol*, RFC 768, <http://www.ietf.org/rfc/rfc0768.txt>

Turning now to Figs. 3A and 3B, block diagrams of embodiments of computer systems 200A and 200B having remote management arrangements are shown, according to

various aspects of the present invention. In Fig. 3A, an ASF south bridge 212 may include integrated ASF, ACPI, and/or Ethernet capabilities for improved remote manageability.

The computer system 200A of Fig. 3A includes a processor 202, a north bridge 204, 5 memory 206, Advanced Graphics Port (AGP) device 208, a PCI bus 210, a PCI connector 211, the ASF south bridge 212, a battery 213, an AT Attachment (ATA) interface 214, an SMBus 215, a USB interface 216, an LPC bus 218, an input/output controller chip (SuperI/O™) 220, extended BIOS memory 222, and, optionally, a crypto-processor 224 and protected storage 230. It is noted that the north bridge 204 and the ASF south bridge 212 may include only a single chip or a plurality of chips in the "chipset." It is also noted that other buses, devices, and/or subsystems may be included in the computer system 200A as desired, e.g. caches, modems, parallel or serial interfaces, SCSI interfaces, etc.

The processor 202 is coupled to the north bridge 204. The north bridge 204 provides 15 an interface between the processor 202, the memory 206, the AGP device 208, and the PCI bus 210. The ASF south bridge 212 provides an interface between the PCI bus 210 and the peripherals, devices, and subsystems coupled to the IDE interface 214, the SMBus 215, the USB interface 216, and the LPC bus 218. The battery 213 is shown coupled to the ASF south bridge 212. The Super I/O™ chip 220, the extended BIOS 222, and the crypto-processor 224 20 are coupled to the LPC bus 218. The protected storage 230 is coupled through the crypto-processor 224.

The north bridge 204 provides communications access between and/or among the processor 202, memory 206, the AGP device 208, devices coupled to the PCI bus 210 and 25 devices and subsystems coupled to the ASF south bridge 212. Typically, removable

5 peripheral devices are inserted into PCI "slots," shown here as the PCI connector 211, that connect to the PCI bus 210 to couple to the computer system 200A. Alternatively, devices located on a motherboard may be directly connected to the PCI bus 210. The SMBus 215 is shown "integrated" with the PCI bus 210 by using pins in the PCI connector 211 for a portion of the SMBus 215 connections.

10 The ASF south bridge 212 provides an interface between the PCI bus 210 and various devices and subsystems, such as a modem, a printer, keyboard, mouse, etc., which are generally coupled to the computer system 200A through the LPC bus 218 (or its predecessors, such as the X-bus or the ISA bus). The ASF south bridge 212 includes logic used to interface the devices to the rest of computer system 200A through the IDE interface 214, the SMBus 215, preferably supporting masters external to the ASF south bridge 212, the USB interface 216, and the LPC bus 218.

15 It is also noted that the operations of the LPC bus 218 may correspond to the prior art Low Pin Count Interface Specification Revision 1.0 of September 29, 1997. The operations of the LPC bus 218 may also correspond to the extended LPC bus disclosed in the LPC Extension Application previously incorporated herein by reference.

20 The extended BIOS 222 includes additional memory locations different from or in addition to those memory locations in the BIOS memory 122. The additional memory locations may have specific read/write permissions and/or be secure memory locations. Additional details may be found in the Secure Execution Mode Applications previously incorporated herein by reference. Memory addressing for the extended BIOS 222 may be as 25 taught in the LPC Extension Application previously incorporated herein by reference. The

crypto-processor 224 may provide security for the protected storage 230. Various embodiments for accessing the protected storage 230 through the crypto-processor 224 are provided in the Secure Execution Mode Applications previously incorporated herein by reference.

5

As mentioned above, the ASF south bridge 212 may include integrated ASF, ACPI, and/or Ethernet functionality, according to various aspects of the present invention. As there is no ASF NIC 109 in the computer system 200A, according to one aspect of the present invention, the ASF south bridge 212 recognizes that it must be a master ASF controller for the computer system 200A, during a power-up cycle. The computer system 200A may advantageously boot faster than the computer system 100 by initiating the ASF and/or ACPI assets in the ASF south bridge 212 during the main portion of the BIOS loading since the ASF, ACPI, and/or Ethernet hardware are known to the BIOS code writer before the BIOS code is written. The BIOS code itself may then be enlarged to include any or all ASF, ACPI, and/or Ethernet initialization data and/or firmware. Additional details of various embodiments of the present invention are given below.

10  
15

15

20

In Fig. 3B, the computer system 200B differs from the computer system 200A in that the computer system 200B includes the ASF NIC 109 at the PCI connector 211. In the computer system 200B, the ASF south bridge 212, according to one aspect of the present invention should recognize that it should be an ASF slave to the ASF NIC 109.

25

The Secure Execution Mode Applications previously incorporated herein by reference teach that power management functions may be performed inside a secure execution mode (SEM), including using security hardware integrated into the south bridge. One current

standard for power management and configuration is the ACPI Specification. According to the ACPI specification, control methods, a type of instruction, tell the computer system to perform an operation. The ACPI specification does not know how to carry out any of the instructions. The ACPI specification only defines the calls, and the software must be written 5 to carry out the calls in a proscribed manner. The proscribed manner of the ACPI specification is very restrictive. One cannot access some registers in your hardware. To access those registers, one can generate an SMI# (System Management Interrupt) to enter SMM and read these registers, as taught in the Secure Execution Mode Applications previously incorporated herein by reference. As power management has the potential to be abused e.g. change the processor voltage and frequency, raised above operating limits to destroy the processor, or lowered below operating limits leading to a denial of service, ACPI calls should be carried out in a secure manner, such as inside SEM.

Inside SEM, each ACPI request can be checked against some internal rules for safe 15 behavior. Using terminology more completely described in the Secure Execution Mode Applications previously incorporated herein by reference, the ACPI request would be placed in an “inbox” (incoming-only memory locations in the south bridge) of a “mailbox” (one-direction-only memory locations in the south bridge), parameter values read from the inbox, the ACPI request evaluated using the inbox parameters for acceptability, and then either 20 fulfill the request or not, based on the evaluation results. For additional details of various embodiments, see the Secure Execution Mode Applications previously incorporated herein by reference, including Figs. 6, 42A, and 42B therein.

System Management Mode (SMM) is a mode of operation in the computer system 25 that was implemented to conserve power. The SMM was created for the fourth generation

x86 processors, and is different from x86 operating mode. As newer x86 generation processors have appeared, the SMM has become relatively transparent to the operating system. That is, computer systems enter and leave the SMM with little or no impact on the operating system.

5

In Fig. 4, one embodiment of the ASF south bridge 212 is illustrated, according to various aspects of the present invention. As shown, an internal south bridge bus 302 couples a south bridge register 304 with an internal bus interface 338 of an Ethernet controller 344 and an LPC bridge 330. The south bridge register 304 also couples to an SMI request register 306, an ASF configuration register 308, a failure recovery timer (WDT) 310, a CPU-MC (microcontroller) interrupt register 312, a CPU-MC data exchange register 314, an ACPI interface 316, an ASF status register 318, and a south bridge register bridge 334. The south bridge register bridge 334 also couples to an MC address/data (A/D) bus 322.

10  
15

Also coupled to the MC A/D bus 322 are a memory 324, an ASF transmit (Tx) buffer 326, an ASF receive (Rx) buffer 328, the LPC bridge 330, an RMCP set command unit 336, and an embedded microcontroller (MC) 320. The MC 320 is also coupled to the WDT 310 and coupled to receive an interrupt (INT) from the CPU-MC interrupt register 312 and the ACPI interface 316. The ACPI interface 316 also generates an SCI interrupt request. The 20 ASF status register 318 also generates an interrupt request. The embedded Ethernet controller also includes a Rx buffer 342 coupled to the ASF Rx buffer 328, a Tx buffer 340 coupled to the ASF Tx buffer 326, and an Ethernet core 344, including a register 346. The Ethernet core 344 is shown coupled to a PHy 348 through an MII (Machine Independent Interface). The PHy 348 may be external to the ASF south bridge 212.

25

The MC 320 couples to the SMBus 215, not shown. The MC 320 may use software-drive I/O ports for the SMBus protocol, according to one aspect of the present invention, using so-called “chapter 13 interfaces” of the ACPI Specification, named from their definition given in chapter 13 of the ACPI Specification. In this embodiment and other 5 embodiments, the processor (CPU) 202 can master the SMBus 215. The MC 320 may store assignable addresses in the memory 324, with fixed motherboard-resident legacy sensor addresses store in the BIOS ROM 122 or the extended BIOS 222. When the ASF NIC 109 is present and the ASF south bridge 212 is operating in slave mode, any sensors internal to the ASF south bridge 212 should be visible to the ASF NIC 109.

10  
15

The embedded Ethernet controller, including the Ethernet core 344, may be configured at boot time from either BIOS code stored in the extended BIOS 222 or by the MC 320 reading values from an EEPROM, not shown, and writing the register 346. It is noted that the register 346 may include a plurality of storage locations or a plurality of registers each with one or more storage locations.

Note that the MC 320 may have some number of general-purpose I/O pins, not shown. The input pins may be used to generate panic interrupts to the MC 320. The output pins may be used to control motherboard 101 functions that are desired when the processor 20 202 may be “hung” and for ASF slave mode panic generation. The ASF slave mode panic generation may substitute for “pushes” of sensor 103 outputs. The general-purpose I/O inputs may generate an interrupt to the MC 320 or be polled by the MC 320, as desired.

Also note that the MC 320 may be configured to manage, control, monitor, and/or 25 provide other functionality for the ASF south bridge 212 besides ASF. Other functionality

may include security, including SEM functionality, system health checking, including ACPI, or other functionality consistent with the teachings herein.

5 The SMI request register 306 is configured to generate an SMI interrupt when an interrupt vector is written to the SMI request register 306. The interrupt vector is passed to an interrupt controller, not shown. It is noted that the SMI request register 306 may be in addition to or the same as the corresponding SMM initiator or SMM initiation register of the Secure Execution Mode Applications previously incorporated herein by reference.

10 The memory 324 may include ROM and/or RAM, as desired. The MC 320 may read configuration data from ROM in the memory 324 and shadow the configuration data in RAM in the memory 324. The configuration data may be stored in the extended BIOS 222 and shadowed in the RAM. Note that the ACPI interface 316 couples to the power/system management core 233, shown in Fig. 3, in the ASF south bridge 212.

15

In one embodiment, the ASF configuration register 308 is a plug and play configuration register for the MC 320 configured for ASF. While ASF is primarily used when the operating system is absent (*e.g.*, not yet loaded at boot time or hung), ASF does interact with the operating system.

20

In one embodiment, the MC 320 is a conventionally available microcontroller, such as an embedded 8051 microcontroller. The 8051 microcontroller and related microcontrollers have well-known functionality in the art. Typical functionality of the 8051 microcontroller includes a central processing unit with a Boolean processor optimized for one-bit operations, 25 five or six interrupts, with two external and two priority levels, two or three timers or

counters, often 16-bit, a programmable full-duplex serial port with data rate defined by one of the timers, 32 I/O lines often as four 8-bit ports, RAM, and optional ROM. The 8051 microcontroller is known to exist in a multitude of varieties, each variation being embraced herein. Other microcontroller and microprocessor designs are also contemplated as the MC 5 320.

Fig. 5 illustrates the RTC battery well 225 of the ASF south bridge 212, according to the present invention. In addition to SB RAM 226, divided into CMOS RAM 226A and RTC RAM 226B, the RTC battery well 225 includes a clock circuit 228, a status register 250, and an enable register 252. The RTC RAM 226B includes checksum data 227 and clock data 229. The battery 213 is coupled to provide power to the contents of the RTC battery well 225. The status register 250 is configured to store status information for the ASF capabilities of the computer system 200. The enable register 252 is configured to store a master bit that, when set, indicates that the ASF NIC 109 is not present. A slave bit may alternatively be stored that, when set, indicates that the ASF NIC 109 is present. It is noted that ASF registers 250 and 252 shown in Fig. 5 may each separately include one or more storage locations or a plurality of registers each having one or more storage locations.

The ASF south bridge 212 also includes, outside the RTC battery well 225, a CPU 20 interface 232, power and system management units 233, and various bus interface logic circuits 234. Time and date data from the clock circuit 228 are stored as the clock data 229 in the RTC RAM 226B. The checksum data 227 in the RTC RAM 226B may be calculated based on the CMOS RAM 226A data and stored by the BIOS code during the boot process. The CPU interface 232 may include interrupt signal controllers and processor signal

controllers. The power and system management units 233 may include an ACPI (Advanced Configuration and Power Interface) controller.

5 Fig. 6 illustrates a flowchart of an embodiment of a method of initializing a computer system including the ASF south bridge 212. Various steps shown in Fig. 2 that are not shown or replaced in Fig. 6 are also contemplated as included in Fig. 6.

During initialization, the processor 202 reads the default jump location. The default jump location in memory is usually at a location such as FFFF0h. The processor 202 performs a jump to the appropriate BIOS code location (e.g. FFFF0h) in the ROM BIOS 222, copies the BIOS code to the RAM memory 206, and begins processing the BIOS code instructions from the RAM memory 206, in block 405. Processing the BIOS code instructions includes checking for the presence of an ASF NIC 109.

15 If the ASF NIC 109 is present, in decision block 410, then the method continues with block 415. If the ASF NIC 109 is not present, in decision block 410, then the method continues with block 420.

20 If the ASF NIC 109 is present, then the ASF south bridge 212 is configured as a slave to the ASF NIC 109, in block 415. If the ASF NIC 109 is not present, then the ASF south bridge 212 is configured as a master ASF device, in block 420. Blocks 415 and 420 are each followed by block 425.

25 The BIOS code, processed by the processor 202, performs a power-on self test (POST), in block 425. The BIOS code next looks for additional BIOS code, such as from a

video controller, IDE controller, SCSI controller, etc. and displays a start-up information screen, in block 430. The BIOS code may perform additional system tests, such as a RAM memory count-up test, and a system inventory, including identifying COM (serial) and LPT (parallel) ports, in block 435. The BIOS code also identifies plug-and-play devices and other 5 similar devices and then displays a summary screen of devices identified, in block 440. The BIOS code identifies the boot location, and the corresponding boot sector, in block 445.

Configuring the ASF south bridge 212 as a slave to the ASF NIC 109, in block 415, may include setting a bit indicating the slave condition in the ASF enable register 252. Configuring the ASF south bridge 212 as the ASF master, in block 420, may include setting a bit indicating the master condition in the ASF enable register 252.

Fig. 7A illustrates a flowchart of an embodiment of a method 500 for operating a computer system including the ASF south bridge 212 in slave mode, according to one aspect 15 of the present invention. The method 500, in one embodiment, may be performed by a slave-mode application that is storable in the computer system 200. In slave mode, the ASF south bridge 212 responds to reads of internal sensor status by the ASF NIC 109, in block 505. The ASF south bridge 212 in slave mode responds to SMBus 215 polls originating on the ASF NIC 109, in block 510. The ASF south bridge 212 in slave mode also provides control points 20 for the ASF NIC 109, allowing the ASF NIC 109 to reset the computer system 200 and cycle the power to the computer system 200.

Fig. 7B illustrates a flowchart of an embodiment of a method 600 for operating a computer system including the ASF south bridge 212 in master mode, according to one 25 aspect of the present invention. The method 600, in one embodiment, may be performed by a

master-mode application that is storable in the computer system 200. In master mode, the ASF south bridge 212 actively polls (exemplary polling engine shown in Fig. 7C) external sensors coupled to the SMBus 215 at a programmable polling rate, in block 605. The ASF south bridge 212 in master mode actively polls or otherwise monitors internal sensor states, 5 in block 610. The ASF south bridge 212 in master mode may generate interrupts and/or respond to interrupts, in block 615. Resulting external sensor status values are combined with internally monitored sensor values and reported to the remote management server 90 via the Ethernet core 344 in the ASF south bridge 212, in block 620.

100-200-300-400-500-600-700-800-900

Fig. 7C illustrates a block diagram of a polling engine 705 in accordance with one embodiment of the present invention. A polling routine 710, based on the sensors listed in a sensor address table 715, builds a current “state of the computer system” table 720. Thus, in one embodiment, the polling engine 705, depending on the mode (master or slave) of execution, continually accesses various internal and external sensors and builds a model of 15 the state of the computer. The polling engine 705 may take action based on inappropriate (or undesirable) sensor values at the end of each polling cycle by sending PET messages, interrupts to the CPU of the computer system 200, and the like. In one embodiment, the CPU of the computer system 200 may access the table 720 directly, as indicated by line 725. Furthermore, in one embodiment, the polling engine 705 uses the values stored in the table 20 720 to construct responses (RMCP frames) to requests from the management console. The polling engine 705, in one embodiment, includes a timer 730 that defines the polling frequency set by a polling rate register (not shown). The timer 730 may be hardware based, or, alternatively, software based.

In the master mode, both internal and external sensors may be read, while in the slave mode, typically only the internal sensors are read. The external sensors 735(1-m) are read via the SMBus interface 740, in one embodiment. The SMBus interface 740 may report alerts over line 745. The polling operation may initiate master cycles on the SMBus 215 to read the 5 various sensors in the computer system 200. As mentioned, the addresses and types of the sensors are stored in the sensor address table 715. An exemplary sensor address table 715 is shown in Fig. 7D.

Referring now to Figure 8, the south bridge 212 comprises a master control loop 805 that includes a queue 810 having a plurality of entries 815(1-n), where the last entry 815(n) in the queue 810 includes a task picker. As more fully explained below, the master control loop 805, which provides three general functions of handling interrupts, scheduling tasks, and idling, may be used by the master controller 320 to execute a variety of tasks. Although the queue 810 is described in the context of the south bridge 212, it should be appreciated that the 15 queue 810 may be employed in a wide variety of other applications where it is desirable to execute one or more independent tasks and where fault-tolerance is desired should one or more of these tasks fail to complete successfully. Furthermore, the master control loop 805 is not limited to ASF applications, and thus may be applicable to non-ASF based applications as well.

20

As mentioned, in the illustrated embodiment, the master control loop 805 includes the task picker in the last entry 815(n) of the queue 810. One or more tasks, when posted in the queue 810, are executed in a preselected order. The exemplary queue 810 illustrated in Figure 8 contains 4 tasks that have been posted in entries 815(1-4). For sake of illustration, it 25 is assumed that task #1 in the entry 815(1) is the oldest task (*i.e.*, was the first posted task),

and that task #4 stored in the entry 815(4) is the newest task to be posted in the queue 810. In accordance with one embodiment, the task picker selects the tasks from the queue 810 based on a priority scheme for execution. For example, in one embodiment, the preselected order may be based on a first-in, first-out scheme (*i.e.*, the order in which the tasks are posted 5 in the queue 810). Thus, in the illustrated example of Figure 9, task #1 (the oldest task) would be executed first and task #4 (the newest task) would be executed last. In other embodiments, the tasks may be executed using a user-designated priority scheme.

For ease of illustration, it is herein assumed that the task picker selects the tasks in the order in which they are posted. That is, in the illustrated embodiment, the task picker selects the oldest task in the queue 810 for execution, as described in more detail below. Once the task picker selects the task for execution, the selected task executes to completion and returns control to the task picker, which then selects the next task in the queue 810 for execution. The tasks are removed from the queue 810 upon execution. When no more tasks remain in the queue 810, the task picker, in one embodiment, continues to execute itself, and thus stays 10 15 in an idle mode, until other tasks are posted in the queue 810.

The south bridge 212, in the illustrated embodiment, includes a failure recovery timer 812 that is capable of generating interrupts that may be detected by the master controller 320. 20 In particular, the failure recovery timer 812, in one embodiment, may generate an interrupt during preselected time intervals. The preselected timer interval may be based, for example, on the amount of time that is required for the tasks in the queue 810 to complete executing. In one embodiment, the failure recovery timer 812, which may be a non-maskable timer that 25 is implemented in hardware, may contain a digital counter that counts down to zero at a constant speed from a preset number. The counter speed may be kept constant by a

conventional clock circuit (not shown). If the counter reaches zero before the task that is currently executing completes, the failure recovery timer 812 generates an interrupt. The failure recovery timer 812 may be resetable. The failure recovery timer 812, in one embodiment, may have a timeout value that is longer than any single task is expected to live.

5 As such, a detection of an interrupt generated by the failure recovery timer 812 may be an indication that one or more tasks may be hung, and thus unable to complete.

In the illustrated embodiment, the south bridge 212 includes a repetitive timer 814 that, as described later in greater detail below, generates interrupts at fixed time intervals to handle tasks that are repetitively invoked.

Referring now to Figure 9, a flow diagram of the master control loop 805 in accordance with one embodiment of the present invention is shown. The flow diagram of Figure 9 begins with the failure recovery timer 812 being reset (at 910). The task picker, stored in the queue 810, is executed (at 920). The task picker determines (at 925) if any tasks (other than the task picker) have been posted in the queue 810. If it is determined (at 925) that no tasks other than the task picker exist in the queue 810, then the failure recovery timer 812 is reset (at 910) and the task picker is once again executed (at 920). The above-described routine may continue to repeat until another task is posted in the queue 810. In this manner, 15 the task picker forms an idle loop for periods during which no new tasks are posted in the queue 810.

20

If it is determined (at 925) that a task other than the task picker exists in the queue 810, the task picker identifies (at 930) the oldest task relative to other tasks (not including the 25 task picker) in the queue 810. If there is only one task in the queue 810, then that task is

identified (at 930) as the oldest task. If there is more than one task in the queue 810, then the task that was posted first in the queue 810 is identified (at 930) as the oldest task. As mentioned earlier, for illustration purposes a first-in, first-out priority scheme is utilized for selecting tasks, although in other embodiments one of a variety of other priority schemes may

5 be employed without departing from the spirit and scope of the instant invention.

The task picker resets (at 940) the failure recovery timer 812, and then passes (at 950) control to the oldest task identified (at 930) and removes the task from the queue 810. It should be appreciated that, in one embodiment, the failure recovery timer 812 may be reset (at 940) at substantially simultaneously the same time control is passed (at 950) to the oldest task identified (at 930) in the queue 810. Additionally, it should be appreciated that, based on design choice, the oldest task may be removed from the queue 810 before, at substantially the same time, or after the control is passed (at 950) to the oldest task that is identified (at 930) in the queue 810.

10  
15

The oldest task identified (at 930) is executed (at 960). It should be appreciated that, in one embodiment, the oldest task may be removed (at 950) from the queue 810 upon execution of that task. Once the oldest task has completed executing (at 960), the failure recovery timer 812 is reset (at 910) and the control is passed to the task picker, which, upon

20 execution (at 920), checks for other tasks in the queue 810. The above-described process is repeated until all of the posted tasks have been executed in the desired order, at which time the task picker stays in an idle mode. In one embodiment, the task picker may periodically poll to determine if any new tasks have been posted in the queue 810.

In one embodiment, the queue 810 may be a pointer that references a starting point of code that performs the desired task. As such, the task picker may look in the queue 810 and determine that a task exists. A “task” may be an address/handle pointing to code to be executed. The task picker may execute that code by making a call, where the argument of 5 that call is the pointer in the queue 810. As a result, a program counter (not shown) is loaded with the address of the first instruction of that task. The task will then execute. The last instruction of that task, upon execution, restores the program counter with the address of the task picker, such that the task picker is executed again upon completion of the previous task.

10 In one embodiment, the master control loop 805 of Figure 9 is capable of responding to interrupts. As more fully explained below, the master control loop 805 may respond to a variety of interrupts, including those that are generated by the failure recovery timer 812 or those that are generated to post one or more tasks in the queue 810. The failure recovery timer 812 may be capable of generating interrupts at preselected time intervals, as indicated 15 above. Generally, these interrupts are utilized by the MC 320 to terminate any tasks that may not have successfully completed within the prescribed time. Interrupts generated by the failure recovery timer 812 are handled by an interrupt service routine 1010, a flow diagram of which is shown in Figure 10, in accordance with one embodiment of the present invention.

20 The interrupt service routine (ISR) 1010 is invoked (at 1020) in response to detecting an interrupt that is generated by the failure recovery timer 812. The ISR 1010, based on detecting the generated interrupt, terminates (at 1025) the task that is currently being executed from the queue 810. Upon termination (at 1025) of the current task, the ISR 1010, in one embodiment, returns the control to the task picker, which may then poll the queue 810 25 for additional, if any, tasks needing to be serviced.

The act of terminating (at 1025) the current task, in one embodiment, may comprise identifying (1032) the task that is currently executing, determining (at 1035) an exit routine associated with that task, and calling (at 1038) the exit routine to terminate. In accordance  
5 with one embodiment, the tasks posted in the queue 810 include an exit routine that may be capable of terminating that task. The act of calling (at 1038) the exit routine, in one embodiment, may include the exit routine setting (at 1045) a “terminate” flag and performing (at 1050) a return from the interrupt. When control is returned from the interrupt to the currently executing task, control returns to the task picker in the queue 810. Thus, the above-  
10 identified blocks 1045 and 1050, in one embodiment, may be performed by the exit routine associated with the currently executing task. The exit routine associated with a particular task that is currently executing may perform, if invoked, additional cleanup steps to facilitate the termination of that task.

15 Referring now to Figure 11, an alternative embodiment of a flow diagram of the master control loop 805 is shown. The illustrated alternative embodiment of the master control loop 805 of Figure 11 is similar to that of Figure 9 except this embodiment does not require the failure recovery timer 812 to be reset. Accordingly, as indicated by the use of like reference numerals, the flow diagram of Figure 11 includes the same blocks as that of Figure  
20 9 but omits blocks 910, 940, which call for resetting the failure recovery timer 812. As described in more detail below, the master control loop 805 of Figure 11, in the illustrated embodiment, does not require the failure recovery timer 812 to be reset before execution of each task in the queue 810.

In one embodiment, the oldest task, when executed (at 960), programs the failure recovery timer 812 to generate an interrupt after a preselected time interval, where the preselected time interval substantially corresponds to the amount of time required for the task to complete execution. Thus, if the task fails to complete executing within the preselected 5 time interval, it may be terminated. In this manner, the failure recovery timer 812 may be programmed to generate an interrupt at different time intervals, depending on the task that is executed at that time. In one embodiment, the preselected time interval may be any desirable time interval greater than the time required for the task to complete executing.

40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65  
66  
67  
68  
69  
70  
71  
72  
73  
74  
75  
76  
77  
78  
79  
80  
81  
82  
83  
84  
85  
86  
87  
88  
89  
90  
91  
92  
93  
94  
95  
96  
97  
98  
99  
100  
101  
102  
103  
104  
105  
106  
107  
108  
109  
110  
111  
112  
113  
114  
115  
116  
117  
118  
119  
120  
121  
122  
123  
124  
125  
126  
127  
128  
129  
130  
131  
132  
133  
134  
135  
136  
137  
138  
139  
140  
141  
142  
143  
144  
145  
146  
147  
148  
149  
150  
151  
152  
153  
154  
155  
156  
157  
158  
159  
160  
161  
162  
163  
164  
165  
166  
167  
168  
169  
170  
171  
172  
173  
174  
175  
176  
177  
178  
179  
180  
181  
182  
183  
184  
185  
186  
187  
188  
189  
190  
191  
192  
193  
194  
195  
196  
197  
198  
199  
200  
201  
202  
203  
204  
205  
206  
207  
208  
209  
210  
211  
212  
213  
214  
215  
216  
217  
218  
219  
220  
221  
222  
223  
224  
225  
226  
227  
228  
229  
230  
231  
232  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
300  
301  
302  
303  
304  
305  
306  
307  
308  
309  
310  
311  
312  
313  
314  
315  
316  
317  
318  
319  
320  
321  
322  
323  
324  
325  
326  
327  
328  
329  
330  
331  
332  
333  
334  
335  
336  
337  
338  
339  
340  
341  
342  
343  
344  
345  
346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377  
378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
401  
402  
403  
404  
405  
406  
407  
408  
409  
410  
411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431  
432  
433  
434  
435  
436  
437  
438  
439  
440  
441  
442  
443  
444  
445  
446  
447  
448  
449  
450  
451  
452  
453  
454  
455  
456  
457  
458  
459  
460  
461  
462  
463  
464  
465  
466  
467  
468  
469  
470  
471  
472  
473  
474  
475  
476  
477  
478  
479  
480  
481  
482  
483  
484  
485  
486  
487  
488  
489  
490  
491  
492  
493  
494  
495  
496  
497  
498  
499  
500  
501  
502  
503  
504  
505  
506  
507  
508  
509  
510  
511  
512  
513  
514  
515  
516  
517  
518  
519  
520  
521  
522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539  
540  
541  
542  
543  
544  
545  
546  
547  
548  
549  
550  
551  
552  
553  
554  
555  
556  
557  
558  
559  
5510  
5511  
5512  
5513  
5514  
5515  
5516  
5517  
5518  
5519  
5520  
5521  
5522  
5523  
5524  
5525  
5526  
5527  
5528  
5529  
5530  
5531  
5532  
5533  
5534  
5535  
5536  
5537  
5538  
5539  
55310  
55311  
55312  
55313  
55314  
55315  
55316  
55317  
55318  
55319  
55320  
55321  
55322  
55323  
55324  
55325  
55326  
55327  
55328  
55329  
55330  
55331  
55332  
55333  
55334  
55335  
55336  
55337  
55338  
55339  
55340  
55341  
55342  
55343  
55344  
55345  
55346  
55347  
55348  
55349  
55350  
55351  
55352  
55353  
55354  
55355  
55356  
55357  
55358  
55359  
55360  
55361  
55362  
55363  
55364  
55365  
55366  
55367  
55368  
55369  
55370  
55371  
55372  
55373  
55374  
55375  
55376  
55377  
55378  
55379  
55380  
55381  
55382  
55383  
55384  
55385  
55386  
55387  
55388  
55389  
55390  
55391  
55392  
55393  
55394  
55395  
55396  
55397  
55398  
55399  
553100  
553101  
553102  
553103  
553104  
553105  
553106  
553107  
553108  
553109  
553110  
553111  
553112  
553113  
553114  
553115  
553116  
553117  
553118  
553119  
553120  
553121  
553122  
553123  
553124  
553125  
553126  
553127  
553128  
553129  
553130  
553131  
553132  
553133  
553134  
553135  
553136  
553137  
553138  
553139  
553140  
553141  
553142  
553143  
553144  
553145  
553146  
553147  
553148  
553149  
553150  
553151  
553152  
553153  
553154  
553155  
553156  
553157  
553158  
553159  
553160  
553161  
553162  
553163  
553164  
553165  
553166  
553167  
553168  
553169  
553170  
553171  
553172  
553173  
553174  
553175  
553176  
553177  
553178  
553179  
553180  
553181  
553182  
553183  
553184  
553185  
553186  
553187  
553188  
553189  
553190  
553191  
553192  
553193  
553194  
553195  
553196  
553197  
553198  
553199  
553200  
553201  
553202  
553203  
553204  
553205  
553206  
553207  
553208  
553209  
553210  
553211  
553212  
553213  
553214  
553215  
553216  
553217  
553218  
553219  
553220  
553221  
553222  
553223  
553224  
553225  
553226  
553227  
553228  
553229  
5532210  
5532211  
5532212  
5532213  
5532214  
5532215  
5532216  
5532217  
5532218  
5532219  
5532220  
5532221  
5532222  
5532223  
5532224  
5532225  
5532226  
5532227  
5532228  
5532229  
55322210  
55322211  
55322212  
55322213  
55322214  
55322215  
55322216  
55322217  
55322218  
55322219  
55322220  
55322221  
55322222  
55322223  
55322224  
55322225  
55322226  
55322227  
55322228  
55322229  
553222210  
553222211  
553222212  
553222213  
553222214  
553222215  
553222216  
553222217  
553222218  
553222219  
553222220  
553222221  
553222222  
553222223  
553222224  
553222225  
553222226  
553222227  
553222228  
553222229  
5532222210  
5532222211  
5532222212  
5532222213  
5532222214  
5532222215  
5532222216  
5532222217  
5532222218  
5532222219  
5532222220  
5532222221  
5532222222  
5532222223  
5532222224  
5532222225  
5532222226  
5532222227  
5532222228  
5532222229  
55322222210  
55322222211  
55322222212  
55322222213  
55322222214  
55322222215  
55322222216  
55322222217  
55322222218  
55322222219  
55322222220  
55322222221  
55322222222  
55322222223  
55322222224  
55322222225  
55322222226  
55322222227  
55322222228  
55322222229  
553222222210  
553222222211  
553222222212  
553222222213  
553222222214  
553222222215  
553222222216  
553222222217  
553222222218  
553222222219  
553222222220  
553222222221  
553222222222  
553222222223  
553222222224  
553222222225  
553222222226  
553222222227  
553222222228  
553222222229  
5532222222210  
5532222222211  
5532222222212  
5532222222213  
5532222222214  
5532222222215  
5532222222216  
5532222222217  
5532222222218  
5532222222219  
5532222222220  
5532222222221  
5532222222222  
5532222222223  
5532222222224  
5532222222225  
5532222222226  
5532222222227  
5532222222228  
5532222222229  
55322222222210  
55322222222211  
55322222222212  
55322222222213  
55322222222214  
55322222222215  
55322222222216  
55322222222217  
55322222222218  
55322222222219  
55322222222220  
55322222222221  
55322222222222  
55322222222223  
55322222222224  
55322222222225  
55322222222226  
55322222222227  
55322222222228  
55322222222229  
553222222222210  
553222222222211  
553222222222212  
553222222222213  
553222222222214  
553222222222215  
553222222222216  
553222222222217  
553222222222218  
553222222222219  
553222222222220  
553222222222221  
553222222222222  
553222222222223  
553222222222224  
553222222222225  
553222222222226  
553222222222227  
553222222222228  
553222222222229  
5532222222222210  
5532222222222211  
5532222222222212  
5532222222222213  
5532222222222214  
5532222222222215  
5532222222222216  
5532222222222217  
5532222222222218  
5532222222222219  
5532222222222220  
5532222222222221  
5532222222222222  
5532222222222223  
5532222222222224  
5532222222222225  
5532222222222226  
5532222222222227  
5532222222222228  
5532222222222229  
55322222222222210  
55322222222222211  
55322222222222212  
55322222222222213  
55322222222222214  
55322222222222215  
55322222222222216  
55322222222222217  
55322222222222218  
55322222222222219  
55322222222222220  
55322222222222221  
55322222222222222  
55322222222222223  
55322222222222224  
55322222222222225  
55322222222222226  
55322222222222227  
55322222222222228  
55322222222222229  
553222222222222210  
553222222222222211  
553222222222222212  
553222222222222213  
553222222222222214  
553222222222222215  
553222222222222216  
553222222222222217  
553222222222222218  
553222222222222219  
553222222222222220  
553222222222222221  
553222222222222222  
553222222222222223  
553222222222222224  
553222222222222225  
553222222222222226  
553222222222222227  
553222222222222228  
553222222222222229  
5532222222222222210  
5532222222222222211  
5532222222222222212  
5532222222222222213  
5532222222222222214  
5532222222222222215  
5532222222222222216  
5532222222222222217  
5532222222222222218  
5532222222222222219  
5532222222222222220  
5532222222222222221  
5532222222222222222  
5532222222222222223  
5532222222222222224  
5532222222222222225  
5532222222222222226  
5532222222222222227  
5532222222222222228  
5532222222222222229  
55322222222222222210  
55322222222222222211  
55322222222222222212  
55322222222222222213  
55322222222222222214  
55322222222222222215  
55322222222222222216  
55322222222222222217  
55322222222222222218  
55322222222222222219  
55322222222222222220  
55322222222222222221  
55322222222222222222  
55322222222222222223  
55322222222222222224  
55322222222222222225  
55322222222222222226  
55322222222222222227  
55322222222222222228  
55322222222222222229  
553222222222222222210  
553222222222222222211  
553222222222222222212  
553222222222222222213  
553222222222222222214  
553222222222222222215  
553222222222222222216  
553222222222222222217  
553222222222222222218  
553222222222222222219  
553222222222222222220  
553222222222222222221  
553222222222222222222  
553222222222222222223  
553222222222222222224  
553222222222222222225  
553222222222222222226  
553222222222222222227  
553222222222222222228  
553222222222222222229  
5532222222222222222210  
5532222222222222222211  
5532222222222222222212  
5532222222222222222213  
5532222222222222222214  
5532222222222222222215  
5532222222222222222216  
5532222222222222222217  
5532222222222222222218  
5532222222222222222219  
5532222222222222222220  
5532222222222222222221  
5532222222222222222222  
5532222222222222222223  
5532222222222222222224  
5532222222222222222225  
5532222222222222222226  
5532222222222222222227  
5532222222222222222228  
5532222222222222222229  
55322222222222222222210  
55322222222222222222211  
55322222222222222222212  
55322222222222222222213  
55322222222222222222214  
55322222222222222222215  
55322222222222222222216  
55322222222222222222217  
55322222222222222222218  
55322222222222222222219  
55322222222222222222220  
55322222222222222222221  
55322222222222222222222  
55322222222222222222223  
55322222222222222222224  
55322222222222222222225  
55322222222222222222226  
55322222222222222222227  
55322222222222222222228  
55322222222222222222229  
553222222222222222222210  
553222222222222222222211  
553222222222222222222212  
553222222222222222222213  
553222222222222222222214  
553222222222222222222215  
553222222222222222222216  
553222222222222222222217  
553222222222222222222218  
553222222222222222222219  
553222222222222222222220  
553222222222222222222221  
553222222222222222222222  
553222222222222222222223  
553222222222222222222224  
553222222222222222222225  
553222222222222222222226  
553222222222222222222227  
553222222222222222222228  
553222222222222222222229  
5532222222222222222222210  
5532222222222222222222211  
5532222222222222222222212  
5532222222222222222222213  
5532222222222222222222214  
5532222222222222222222215  
5532222222222222222222216  
5532222222222222222222217  
5532222222222222222222218  
5532222222222222222222219  
5532222222222222222222220  
5532222222222222222222221  
5532222222222222222222222  
5532222222222222222222223  
5532222222222222222222224  
5532222222222222222222225  
5532222222222222222222226  
5

executing task. If the two task IDs are not the same, then it is an indication that the same task has not been executing between the last two successive interrupts, which means that the task that was logged (at 1215) has since completed successfully and that a different task is in the process of being executed. The “task ID,” in one embodiment, may be a 16-bit sequence 5 number that is incremented each time the task is invoked to reduce the chances of the task that repeats frequently from being mistakenly terminated.

If the ISR 1205 determines (at 1230) that the current task ID is the same as the task ID logged (at 1215), then the ISR 1205 terminates (at 1240) the current task. The current task is terminated (at 1240) because the same task has been executing between two successive interrupts, which may be an indication that the current task is hung or unable to complete execution, considering the fact that under normal conditions the task should have completed within one full interrupt interval. Upon termination (at 1240) of the current task, in one embodiment, control is returned (at 1245) to the task picker.

10  
15

The master control loop 805 of Figure 8 allows the computer system 200 to be more fault tolerant. Fault tolerance is a desirable feature in a system, particularly in a system where polling occurs. In the above-described master control loop 805, errant tasks may be terminated, thereby allowing the computer system 200 to continue operating. In some 20 instances, errant tasks may be a result of an application that continues to post tasks periodically in the queue 810 but the posted tasks fail to complete for selected reasons. For example, a keyboard application may use the queue 810 to poll some hardware in the computer system 200 every few seconds to check if that hardware is present. If, however, the polled hardware is, for instance, temporarily removed from the computer system 200, the 25 tasks checking for such hardware may fail, and thus not complete. The master control loop

805, in one embodiment, continues to terminate the hung tasks until the hardware is once again detected in the computer system 200. In this manner, the computer system 200 operates in a substantially seamless manner, at least from the end user's perspective, from the point the polled hardware is not detected to a point the hardware is eventually detected in the 5 computer system 200.

A variety of applications or devices in the computer system 200 (see Figures 3A and 3B) may post tasks in the queue 810 for execution. For example, the CPU of the computer system 200 may post tasks in the queue 810 to access one or more sensors coupled to the SMBus 215 (see Figures 3A and 3B). Once the task is posted in the queue 810, the MC 320 may execute the task to provide the requested information from the sensors on the SMBus 215. In one embodiment, the CPU of the computer system 200 may generate an MC 320 interrupt request and pass a task vector to the MC 320. The CPU may write to an interrupt trap register of the MC 320 that generates an MC 320 interrupt. The write operation may 10 generate an interrupt, and thereafter data may be exchanged via a data exchange register, for example. 15

In addition to the CPU, other devices and applications, such as the master-mode application, slave-mode application, keyboard application (*i.e.*, application for managing 20 keyboard operation), and the like, may also post a variety of tasks in the queue 810 for execution. Some exemplary tasks that may be posted in the queue 810, and the manner in which these tasks may be posted, are described in greater detail below.

When configured to operate in the master mode, the master-mode application of the 25 south bridge 212 actively polls external sensors coupled to the SMBus 215 as well as internal

sensors. The term “sensor,” as utilized herein, refers to any hardware source of status information. “Polling” the external or internal sensors periodically is one example of a task that may be posted in the queue 810 for execution. Polling is one example of a task that is repetitively invoked. Repetitive tasks, in accordance with one embodiment, may be handled 5 in a variety of ways, including through the use of interrupt service routines and posting the task on the queue 810 in response to an interrupt generated by the repetitive timer 814. For performance and reliability reasons, it may be desirable not to burden interrupt service routines to perform time-consuming tasks. As such, in some instances it may be desirable to use interrupt service routines generally for short operations, such as resetting a timer or 10 reading a status register to detect a change, for example.

100-105

15

20

25

The repetitive timer 814, in one embodiment, generates an interrupt at preselected time intervals. The interrupt is in turn serviced by an interrupt service routine that determines the source of the interrupt, determines what is needed to respond to the interrupt, and posts one or more tasks in the queue 810 to properly address or service the interrupt. Once the tasks are posted in the queue 810, the task picker of the master control loop 810 processes those tasks whenever possible. Deferring the tasks to the queue 805 allows the interrupt service routines to make a quick and clean exit. It should be noted that in accordance with one embodiment of the present invention, repetitive (e.g., polling) tasks may be posted by either the master-mode application, the slave-mode application, or any other application requiring such tasks to be serviced.

26

27

28

29

In one embodiment, the repetitive timer 814 may be used to load more than one task into the queue 810 per timeout. In one embodiment, a more robust polling mechanism may be utilized where the interrupt service routine builds one or more programmable timers that

control separate tasks with different frequencies (the repeat time may be an integer multiple of the time base of the repetitive timer 814).

In the master mode, as mentioned above, Ethernet packets are constructed by the 5 master controller 320 and stored in the transmit buffer 326 (see Figure 4). In particular, in one embodiment, various PET, RMCP response, and RMCP ACK/NACK packets are created by the master controller 320 and stored in the transmit buffer 326. Additionally, during the master mode, a variety of packets are extracted from Ethernet packets stored from the receive buffer 328 (see Figure 4). The task of constructing and deconstructing Ethernet frames may be placed on the queue 810, in one embodiment, by the master-mode application. The master controller 320, when possible, may process the tasks from the queue 810 once they are posted.

The south bridge 212, in one embodiment, supports SMBus master emulation and 15 slave emulation modes. In the master mode, the transactions may be initiated by either the master controller 320 or the CPU of the computer system 200. When in the slave mode, the SMBus 215 is the target of messages from the SMBus master on the NIC 109. In the slave mode, the south bridge 212 should recognize transactions targeting its address and respond accordingly and also respond to the inbound address resolution protocol enumeration cycles, 20 in one embodiment. The tasks associated with the SMBus emulation modes may be placed on the queue 810 for execution by the master controller 320.

For the purposes of this disclosure, references to ROM are to be construed as also 25 applying to flash memory and other substantially non-volatile memory types. Note that while the methods of the present invention disclosed herein have been illustrated as flowcharts,

various elements of the flowcharts may be omitted or performed in different order in various embodiments. Note also that the methods of the present invention disclosed herein admit to variations in implementation.

5 Some aspects of the invention as disclosed above may be implemented in hardware or software. Thus, some portions of the detailed descriptions herein are consequently presented in terms of a hardware implemented process and some portions of the detailed descriptions herein are consequently presented in terms of a software-implemented process involving symbolic representations of operations on data bits within a memory of a computing system or computing device. These descriptions and representations are the means used by those in the art to convey most effectively the substance of their work to others skilled in the art using both hardware and software. The process and operation of both require physical manipulations of physical quantities. In software, usually, though not necessarily, these quantities take the form of electrical, magnetic, or optical signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

10

15

It should be borne in mind, however, that all of these and similar terms are to be 20 associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated or otherwise as may be apparent, throughout the present disclosure, these descriptions refer to the action and processes of an electronic device, that manipulates and transforms data represented as physical (electronic, magnetic, or optical) quantities within some electronic device's storage into other data similarly 25 represented as physical quantities within the storage, or in transmission or display devices.

Exemplary of the terms denoting such a description are, without limitation, the terms "processing," "computing," "calculating," "determining," "displaying," and the like.

Note also that the software-implemented aspects of the invention are typically 5 encoded on some form of program storage medium or implemented over some type of transmission medium. The program storage medium may be magnetic (e.g., a floppy disk or a hard drive) or optical (e.g., a compact disk read only memory, or "CD ROM"), and may be read only or random access. Similarly, the transmission medium may be twisted wire pairs, coaxial cable, optical fiber, or some other suitable transmission medium known to the art. 10 The invention is not limited by these aspects of any given implementation.

The particular embodiments disclosed above are illustrative only, as the invention 15 may be modified and practiced in different but equivalent manners apparent to those skilled in the art having the benefit of the teachings herein. Furthermore, no limitations are intended to the details of construction or design herein shown, other than as described in the claims below. It is therefore evident that the particular embodiments disclosed above may be altered or modified and all such variations are considered within the scope and spirit of the invention. Accordingly, the protection sought herein is as set forth in the claims below.